

Articulated Pose Estimation Using Hierarchical Exemplar-Based Models

Jiongxin Liu, Yinxiao Li, Peter Allen, Peter Belhumeur

Columbia University in the City of New York
{liujx09, yli, allen, belhumeur}@cs.columbia.edu

Abstract

Exemplar-based models have achieved great success on localizing the parts of semi-rigid objects. However, their efficacy on highly articulated objects such as humans is yet to be explored. Inspired by hierarchical object representation and recent application of Deep Convolutional Neural Networks (DCNNs) on human pose estimation, we propose a novel formulation that incorporates both hierarchical exemplar-based models and DCNNs in the spatial terms. Specifically, we obtain more expressive spatial models by assuming independence between exemplars at different levels in the hierarchy; we also obtain stronger spatial constraints by inferring the spatial relations between parts at the same level. As our method strikes a good balance between expressiveness and strength of spatial models, it is both effective and generalizable, achieving state-of-the-art results on different benchmarks: Leeds Sports Dataset and CUB-200-2011.

Introduction

Articulated pose estimation from a static image remains a challenge in computer vision. The difficulty lies in the wide variations in the appearance of object due to articulated deformations. Therefore, an effective method generally relies on strong appearance models and expressive spatial models to capture the variations. More importantly, these models should be incorporated into a sensible framework where correct poses do enjoy relatively high likelihood.

There has been great progress in developing appearance and spatial models. Histogram of Gradient (HOG) (Dalal and Triggs 2005) was widely used as the part descriptor. However, HOG is rather weak and introduces visual ambiguities (Vondrick et al. 2013). Recently, Deep Convolutional Neural Networks (DCNNs) have demonstrated excellent performance in building appearance models for object detection (Szegedy, Toshev, and Erhan 2013; Sermanet et al. 2014; Girshick et al. 2014) and human pose estimation (Toshev and Szegedy 2014; Jain et al. 2014; Chen and Yuille 2014; Tompson et al. 2014). Compared with a shallow classifier with hand-crafted features, DCNNs have the capacity of learning more discriminative features.

As for the spatial models, tree-structured model wins popularity by enjoying a simplified representation of object shape. In this model, parts are connected with tree edges modeled as elastic springs. On top of such model, various

formulations have been proposed. (Felzenszwalb and Huttenlocher 2005) designed pictorial structure to combine appearance and spatial terms in a generative way. Discriminatively trained method is more powerful, and easily adapts to extended deformable part models (DPMs) with mixture of parts (Yang and Ramanan 2011) and hierarchical representations (Sun and Savarese 2011; Branson, Perona, and Belongie 2011; Sapp and Taskar 2013; Wang and Li 2013). There are efforts to go beyond the tree structure: (Wang, Tran, and Liao 2011) proposes a loopy model that contains composite parts implemented by Poselets (Bourdev and Malik 2009; Bourdev et al. 2010); (Ramakrishna et al. 2014; Tompson et al. 2014) treat all the other parts as the neighbors of a target part; Grammar model has also been proposed to generalize the tree structure (Rothrock, Park, and Zhu 2013).

As nonparametric spatial models, exemplars are also effective in part localization problems (Sapp, Jordan, and Taskar 2010; Belhumeur et al. 2011; Liu and Belhumeur 2013), as an ensemble of labeled samples (i.e., exemplars) literally capture plausible part configurations without assuming the distribution of part relations. However, the exemplar-based models have not shown much efficacy in human pose estimation due to degraded expressiveness of limited training samples (Liu, Li, and Belhumeur 2014).

In our work, we improve the expressiveness of exemplar-based models by leveraging hierarchy and composite parts: each composite part is treated as an object, with exemplars dictating the configuration of its child parts. By doing so, we obtain exemplar-based models covering a spectrum of granularity. In addition, we propose a novel formulation that captures the interactions between object parts in two aspects: spatial relations between parts at the same granularity level are inferred by DCNNs (inspired by (Chen and Yuille 2014)), and constrained by exemplar-based models at that level; spatial relations between parts at different levels follow the parent-child relations in the hierarchy, which are well maintained in the bottom-up inference through exemplars. These efforts together allow us to only model the part relations in each layer via DCNNs without impairing the strength of the method. In some sense, our formulation is tailored for exemplar-based inference, which differs from other hierarchical models such as multi-layer DPMs. Also note that we use grouped parts to optimize the individual parts jointly, which differs from (Bourdev and Malik 2009).

The Approach

Our method features a hierarchical representation of object. We will first describe the relevant notations and introduce hierarchical exemplars. Then we will explain our formulation of pose estimation. In the end, a comparison with relevant techniques will be addressed.

Hierarchical Exemplars

A hierarchical object (exemplar) contains two types of parts: *atomic part* and *composite part*. An atomic part i is at the finest level of granularity, and can be annotated as a keypoint with pixel location x_i (e.g., elbow). A composite part k is the composite of its child parts (e.g., arm = {shoulder, elbow, wrist}), and is denoted as a tight bounding box b_k containing all the child parts inside. In our work, a *part configuration* X is denoted as the locations of atomic parts $[x_1, \dots, x_N]$ where N is the total number of atomic parts.

Now, we define the *spatial relation* between parts of the same type. For atomic parts i and j , their offset $x_j - x_i$ characterizes the relation $r_{i,j}$ (e.g., shoulder is 20 pixels above the elbow). For composite parts k and h , we first assign anchor points a_k and a_h to them. Anchor points are manually determined such that they are relatively rigid w.r.t the articulated deformation. Then we represent the relation $r_{k,h}$ as $[tl(b_h) - a_k, br(b_h) - a_k]$, where $tl(\cdot)$ and $br(\cdot)$ are the top-left and bottom-right corners of part bounding box (please see Fig. 1(a)).

The hierarchical representation follows a tree structure, as shown in Fig. 1(a). Each leaf node denotes an atomic part. A node k at level $l > 1$ corresponds to a composite part $k^{(l)}$ – the union of its children at level $l - 1$ which are denoted as $C(k^{(l)})$. The degree of the tree is not bounded, and the structure of the tree depends on the particular object category. A general rule is: geometrically neighboring parts (at the same level) can form a part at an upper level if their spatial relations can be captured by the training data. Fig. 1(b) shows the instantiation of hierarchical representation for human and bird. As the bird body is relatively more rigid than the human body, the degrees of bird’s internal nodes can be larger, resulting in fewer levels.

The exemplars in previous works (Belhumeur et al. 2011; Liu and Belhumeur 2013) correspond to a depth-2 hierarchy, where all the atomic parts are the children of the unique composite part (i.e., root part). As a result, each exemplar models the relations between all the atomic parts, making the ensemble of training exemplars not capable of capturing unseen poses. Our hierarchical exemplars, however, adapt to a hierarchy with larger depth which gives us multiple composite parts. By treating each composite part as a standalone object, we have exemplars only modeling the spatial relations between its child parts (which are referred to as part pose). We use $\mathcal{M} = \{\mathcal{M}_k^{(l)}\}_{l>1}$ to denote the set of exemplar-based models for all the composite parts, where k is the index, and l denotes the level. By design, the hierarchical exemplars cover a spectrum of granularity with proper decomposition of the object, which dramatically improves the expressiveness of exemplars. Note that the depth had better not go too large, as we still want to make use of

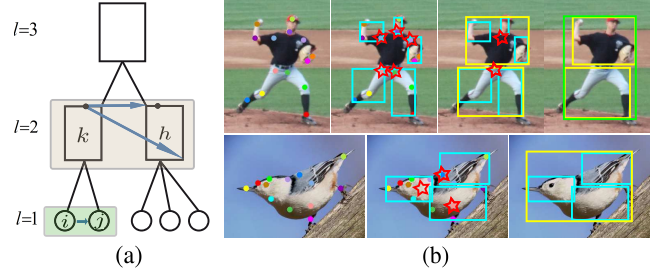


Figure 1: (a) shows the tree-structured hierarchy. The arrows indicate the way of estimating spatial relations (i.e., $r_{i,j}$, $r_{k,h}$) between sibling parts. The black dots on level 2 denote the anchor points. (b) The instantiations of part hierarchy on human and bird. The part levels increase from the left to the right. Each figure shows the parts at the same level with the same color (except for the atomic parts). Immediate children are also plotted for level 2 and above. The stars mark the anchor points for the composite parts.

the strength of exemplars in constraining the configurations of more than two parts.

Formulation

We define an energy function to score a configuration X given an image I and the spatial models \mathcal{M} :

$$S(X|I, \mathcal{M}) = U(X|I) + R(X|I, \mathcal{M}) + w_0, \quad (1)$$

where $U(\cdot)$ is the appearance term, $R(\cdot)$ is the spatial term, and w_0 is the bias parameter. Our goal is to find the best configuration $X^* = \arg \max_X S(X|I, \mathcal{M})$.

Appearance Terms: $U(X|I)$ is a weighted combination of the detection scores for each atomic part:

$$U(X|I) = \sum_{i=1}^N w_i \varphi \left(i|I \left(x_i, s^{(1)}(X) \right) \right), \quad (2)$$

where w_i is the weight parameter, $\varphi(\cdot)$ scores the presence of part i at the location x_i based on the local image patch (Eq. 9), and $s^{(1)}(X)$ denotes the level-1 scaling factor based on X ’s size. As we use sliding-window paradigm for detection, the local image patch is expected to fit the object size and the part’s level.

Spatial Terms: We design multi-level spatial terms to evaluate the part relations. Assuming there are L levels in the object hierarchy, and there are n_l parts at the l -th level, then $R(X|I, \mathcal{M})$ is defined as

$$R(X|I, \mathcal{M}) = \sum_{l=2}^L \sum_{k=1}^{n_l} \Psi \left(p_k^{(l)} | b_k^{(l)}, I, \mathcal{M}_k^{(l)} \right), \quad (3)$$

where $b_k^{(l)}$ denotes the bounding box of part $k^{(l)}$, $p_k^{(l)}$ denotes the pose of $k^{(l)}$, $\mathcal{M}_k^{(l)}$ denotes the corresponding spatial models, and $\Psi(\cdot)$ scores $p_k^{(l)}$ based on both appearance and spatial models. Note that $p_k^{(l)}$ is defined as the spatial relations between the children of $k^{(l)}$.

We now elaborate the derivation of $\Psi(\cdot)$. Using exemplar-based models, we can assume $\mathcal{M}_k^{(l)}$ contains T exemplars $\{X_i\}_{i=1,\dots,T}$, each of which dictates a particular pose p_i (e.g., an example of raised arm). Here, we drop the subscript k and superscript l for clarity. With these in hand, we evaluate $\Psi(\cdot)$ as the combination of two terms:

$$\Psi(p_k^{(l)}|b_k^{(l)}, I, \mathcal{M}_k^{(l)}) = \alpha_k^{(l)} \phi(p_o|I(b_k^{(l)}, s^{(l-1)}(X))) + \beta_k^{(l)} \psi(p_k^{(l)}, p_o), \quad (4)$$

where $\alpha_k^{(l)}$ and $\beta_k^{(l)}$ are the weight parameters, p_o (corresponding to exemplar X_o) is the pose that best fits $p_k^{(l)}$. $\phi(\cdot)$ evaluates the likelihood of pose p_o being present in the image region at $b_k^{(l)}$ (Eq. 10, 11). The relevant image patches also need to be resized as Eq. 2. $\psi(\cdot)$ measures the similarity between $p_k^{(l)}$ and p_o as

$$\psi(p_k^{(l)}, p_o) = -\min_t \|\vec{X}_k^{(l)} - t(\vec{X}_o)\|^2, \quad (5)$$

where t denotes the operation of similarity transformation (the rotation angle is constrained), $\vec{X}_k^{(l)}$ denotes the vectorized locations of parts $C(k^{(l)})$ in X , and \vec{X}_o denotes the vectorized X_o .

As multi-scale image cues and multi-level spatial models are both involved, $\Psi(\cdot)$ covers part relations at different levels of granularity. For instance, at a fine scale (small l), it evaluates whether the arm is folded; at a coarse scale (large l), it evaluates whether the person is seated.

Discussion

We make independence assumption on the spatial models in Eq. 3, which can benefit articulated pose estimation. The reason lies in that it gives us a collection of spatial models that can better handle rare poses. For instance, our models allow a person to exhibit arbitrarily plausible poses at either arm as long as the spatial relations between the two arms are plausible. With such assumption, our formulation still captures the part relations thoroughly and precisely: the relations between sibling parts are encoded explicitly in Eq. 4; the relations between parent and child parts are implicitly enforced (the same X is referred to across the levels).

Below, we address the differences between our method and relevant techniques, such as image dependent spatial relations (Chen and Yuille 2014; Sapp and Taskar 2013; Pishchulin et al. 2013a) and hierarchical models (Sun and Savarese 2011; Wang, Tran, and Liao 2011; Wang and Li 2013):

- Unlike (Chen and Yuille 2014; Sapp and Taskar 2013), our method infers from the image the spatial relations between atomic parts (e.g., elbow and shoulder), as well as the relations between composite parts (e.g., upper body and lower body).
- Unlike (Sapp and Taskar 2013; Pishchulin et al. 2013a), we do not conduct the selection of spatial models upfront as errors in this step are hard to correct afterwards. Instead, our selection of spatial models is based on the configuration under evaluation (the second term in Eq. 4), which avoids pruning the model space too aggressively.

- Unlike (Sun and Savarese 2011; Wang, Tran, and Liao 2011; Wang and Li 2013), our method directly optimizes on the atomic part locations, avoiding the interference from localizing the composite parts. Also, we turn to exemplars to constrain the part relations, rather than using piece-wisely stitched “spring models”.

Inference

The optimization of Eq. 1 does not conform to general message passing framework due to the dependency of p_o on X (Eq. 4) and the interactions between variables x_i across multiple levels (Eq. 3). Therefore, we propose an algorithm (Algorithm 1) which simplifies the evaluation of Eq. 3. Although being approximate, the algorithm is efficient and yields good results. In the following sections, we explain the two major components of the algorithm.

Hypothesize and Test

The first component is *Hypothesize and Test*, which leverages a RANSAC-like procedure of exemplar matching. For this, we rewrite Eq. 3 in a recursive form which scores the subtree rooted at $b_k^{(l)}$ ($l \geq 2$):

$$f(b_k^{(l)}) = \sum_{j \in C(k^{(l)})} f(b_j^{(l-1)}) + \Psi(p_k^{(l)}|b_k^{(l)}, I, \mathcal{M}_k^{(l)}). \quad (6)$$

Note that $f(b_k^{(1)}) = 0$ for any k . By comparing Eq. 6 with Eq. 3, we can see that $f(b_1^{(L)}) = R(X|I, \mathcal{M})$.

Hypothesize and Test is conducted in a bottom-up manner: (1) Given the hypothesized locations of all the parts at level $l-1$ (each part has multiple hypotheses), transform the exemplars at level l to the test image with similarity transformation such that each exemplar’s child parts aligns with two randomly selected hypotheses of atomic parts (if $l=2$), or up to two hypotheses of composite parts (if $l>2$). (2) The geometrically well-aligned exemplars generate hypotheses for the parts at level l . Each hypothesis carries from exemplar the object size, the corresponding subtree, as well as the pose p_o for each node in the subtree. (3) Augment the hypotheses of $k^{(l)}$ (if $l>2$) by swapping their subtrees with geometrically compatible hypotheses at level $l-1$. (4) Evaluate all the hypotheses at level l using Eq. 6 and keep the top-scoring ones. (5) Increment l and go to step (1). Fig. 2 shows examples of the first three steps.

Backtrack

The second component of the algorithm is *Backtrack*. Assuming we have a hypothesis of the root part $b_1^{(L)}$, we can trace down the hierarchy, which gives us p_o in Eq. 4 for each composite part, as well as the hypothesized locations for the atomic parts $\hat{X} = [\hat{x}_1, \hat{x}_2, \dots, \hat{x}_N]$.

The next step is to re-score \hat{X} by obtaining its refined configuration \hat{X}^* . For this purpose, we define $g(\cdot)$ to approximate $S(\cdot)$ in Eq. 1:

$$g(\hat{X}|I, b_1^{(L)}) = U(\hat{X}|I) + \sum_k^{n_2} \beta_k^{(2)} \psi(p_k^{(2)}, p_o) + D, \quad (7)$$

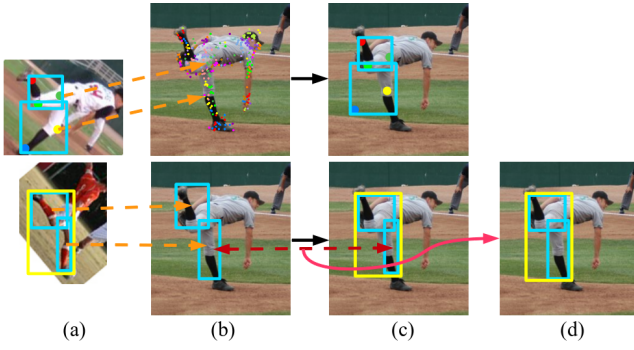


Figure 2: Generate hypotheses from exemplars. The first row corresponds to $l = 2$ and the second row corresponds to $l = 3$. (a) Two training exemplars. (b) The test image overlaid with hypotheses at level $l - 1$. (c) Part hypotheses at level l which are inherited from the exemplars. (d) Augmented hypothesis after swapping hypotheses of child part.

Algorithm 1: Inference Procedure for Pose Estimation

Input: Multi-level exemplars $\{\mathcal{M}^{(l)}\}_{l=2,\dots,L}$;
Multi-level appearance models $\{\mathcal{C}^{(l)}\}_{l=1,\dots,L-1}$;
Test image I ;
Maximum number of hypotheses per part Z ;
Output: The optimal configuration X^* ;
 $hypo^{(1)} \leftarrow$ top Z local maximas from $\mathcal{C}^{(1)}(I)$, $l \leftarrow 2$;
while $l \leq L$ **do**
 $hypo^{(l)} \leftarrow$ randomly align $\mathcal{M}^{(l)}$ with $hypo^{(l-1)}$;
 Augment $hypo^{(l)}$ if $l > 2$;
 Evaluate $hypo^{(l)}$ using Eq. 6 and $\mathcal{C}^{(l-1)}(I)$;
 $hypo^{(l)} \leftarrow$ top-scoring Z $hypo^{(l)}$;
 $l \leftarrow l + 1$;
Refine and re-score $hypo^{(L)}$ through *backtrack*;
 $X^* \leftarrow$ highest-scoring \hat{X}^* ;
return X^* ;

where $D = f(b_1^{(L)}) + w_0$. Such approximation assumes $s(\hat{X})$, p_o and $b_k^{(l)}$ for $l > 2$ change little during the refinement. After plugging Eq. 2 and Eq. 5 into Eq. 7, we can solve each atomic part independently as

$$\hat{x}_i^* = \arg \max_{x_i \in \mathcal{R}(\hat{x}_i)} w_i \varphi(i|I(x_i, s^{(1)}(\hat{X}))) + \beta_k^{(2)} \|x_i - \hat{x}_i\|^2, \quad (8)$$

where part $k^{(2)}$ is the parent of part i , $\mathcal{R}(\hat{x}_i)$ denotes the search region of part i . We define the search region as a circle with radius equal to 15% of the average side length of $b_1^{(L)}$. We evaluate Eq. 8 for all the pixel locations inside the circle, which gives us the highest-scoring location. In the end, we obtain the refined configuration $\hat{X}^* = [\hat{x}_1^*, \hat{x}_2^*, \dots, \hat{x}_N^*]$ with updated score $g(\hat{X}^*|I, b_1^{(L)})$.

Model Learning

In this section, we describe how we learn the appearance models in Eq. 1 (i.e., $\varphi(\cdot)$ and $\phi(\cdot)$), as well as how we learn the weight parameters \mathbf{w} (i.e., w_i , $\alpha_k^{(l)}$, and $\beta_k^{(l)}$).




Figure 3: TOP: The architecture of DCNN-based atomic part detector. It consists of five convolutional layers, two max-pooling layers and three fully-connected layers. The output dimension is $|S|$. BOTTOM: The architecture of DCNN-based models for composite parts. It consists of five convolutional layers, three max-pooling layers and three fully-connected layers. The last two layers branch out, with each branch targeting the possible spatial relations of one composite part to its predefined reference part.

Relations Between Atomic Parts

We follow the method of (Chen and Yuille 2014) to infer the spatial relations between atomic parts. Specifically, we design a DCNN-based multi-class classifier using Caffe (Jia et al. 2014). The architecture is shown in the first row of Fig. 3. Each value in the output corresponds to $p(i, m_{i,j}|I(x, s^{(1)}(X)))$, which is the likelihood of seeing an atomic part i with a certain spatial relation (type $m_{i,j}$) to its predefined neighbor j , at location x . If $i = 0$, then $m_{i,j} \in \{0\}$, indicating the background; if $i \in \{1, \dots, N\}$, then $m_{i,j} \in \{1, \dots, T_{i,j}\}$. By marginalization, we can derive $\varphi(\cdot)$ and $\phi(\cdot)$ as

$$\varphi(i|I(x, s)) = \log(p(i|I(x, s))). \quad (9)$$

$$\phi(p_o|I(b_k^{(2)}, s)) = \sum_{i \in C(k^{(2)})} \log(p(m_{i,j}|i, I(x_i, s))). \quad (10)$$

Note that superscript (1) and X are dropped for clarity, i and j are siblings. To define type $m_{i,j}$, during training, we discretize the orientations of $r_{i,j}$ into $T_{i,j}$ (e.g., 12) uniform bins, and $m_{i,j}$ indicates a particular bin. The training samples are then labeled as $(i, m_{i,j})$.

Relations Between Composite Parts

We build another DCNN-based model to infer the spatial relations between composite parts, as shown in the second row of Fig. 3, the architecture differs from that for atomic parts in multiple aspects. First, as the model targets composite parts which have coarser levels of granularity, the network has a larger receptive field. Second, as there are relatively fewer composite parts than atomic parts, we let all the composite parts share the features in the first several layers (the input patches of different composite parts have different receptive fields). Third, as the composite parts have different granularity with possibly significant overlap with each other, the DCNN branches out to handle them separately.

Assuming the i -th branch corresponds to part i at level $l - 1$ (Note that $l > 2$), then the branch has $|\mathcal{S}_i|$ -dim output with each value being $p(m_{i,j}|i, I(a_i, s^{(l-1)}(X)))$ based on the image patch centered at the anchor point a_i . Assuming the parent of part i is part $k^{(l)}$, then $\phi(\cdot)$ is evaluated as

$$\phi(p_o|I(b_k^{(l)}, s)) = \sum_{i \in C(k^{(l)})} \log(p(m_{i,j}|i, I(b_i, s))). \quad (11)$$

Note that superscript $(l-1)$ and X are dropped for clarity. To train this model, we cluster the relation vector $r_{i,j}$ into $T_{i,j}$ (e.g., 24) clusters (types) for part i , and the training samples are labeled accordingly.

Weight Parameters

Eq. 1 can be written as a dot product $\langle \mathbf{w}, \Phi(X, I, \mathcal{M}) \rangle$. Given a training sample (X, I) , we compute $\Phi(X, I, \mathcal{M})$ as its feature. Each training sample also has a binary label, indicating if the configuration X is correct. Therefore, we build a binary max-margin classifier (Tschantz et al. 2004) to estimate \mathbf{w} , with non-negativity constraints imposed. To avoid over-fitting, the training is conducted on a held-out validation set that was not used to train the DCNNs.

Before training, we augment the positive samples by randomly perturbing their part locations as long as they are reasonably close to the ground-truth locations. To generate the negative samples, we randomly place the configurations of positive samples at the incorrect regions of the training images, with Gaussian noise added to the part locations.




Figure 4: Qualitative results of human pose estimation on LSP dataset (OC annotations). Failures are denoted with red frames, which are due to extreme self-occlusion.

Experiments

We evaluate our method extensively on multiple benchmarks, and conduct diagnostic experiments to show the effect of different components in our method.

LSP Dataset (OC Annotations)

The Leeds Sports Pose (LSP) dataset (Johnson and Everingham 2010) includes 1,000 images for training and 1,000 images for testing, where each image is annotated with 14 joint locations. We augment the training data by left-right

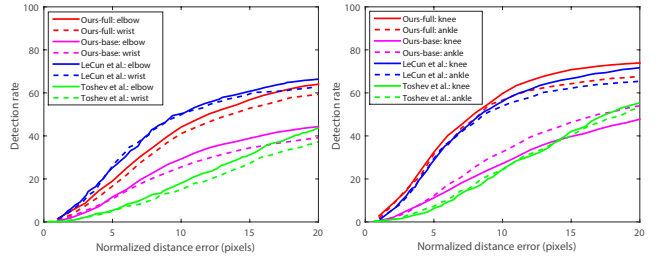


Figure 5: Detection rate vs. normalized error curves on the LSP Extended dataset (PC Annotations). LEFT: arm (elbow and wrist) detection. RIGHT: leg (knee and ankle) detection.

flipping, and rotating through 360° . We use observer-centric (OC) annotations to have fair comparisons with the majority of existing methods. To measure the performance, we use Percentage of Correct Parts (PCP). In PCP measure, a “part” is defined as a line segment connecting two neighboring joints. If both of the segment endpoints (joints) lie within 50% of the length of the ground-truth annotated endpoints, then the part is correct.

In this experiment, we build a hierarchy of four levels for human body. The first level contains the atomic body joints; the second level has five composite parts (Head, Right arm, Left arm, Right leg, and Left leg); the third level has two composite parts (Head&Arms and Legs); and the fourth level corresponds to the whole body. To gain an understanding of the effect of the components of our inference algorithm, we evaluate our full method (which will be referred to as “Ours-full”), and variants of our method (which will be referred to as “Ours-partial”, and “Ours-no-HIER”). Ours-full corresponds to the whole inference algorithm; Ours-partial only conducts the first part of the inference algorithm, traces down the best root hypothesis based on Eq. 6, and outputs the locations of its atomic parts; Ours-no-HIER only uses full-body exemplars (after augmentation) as the spatial models.

The quantitative results of our method as well as its counterparts are listed in Tab. 1. Ours-full generally outperforms the state-of-the-art methods on all the parts. The improvement over IDPR (Chen and Yuille 2014) demonstrates the effect of reasoning multi-level spatial relations. We expect to see even larger improvement if we augment the annotations with midway points between joints. We also experiment with person-centric (PC) annotations on the same dataset, where the accuracy drops slightly. Ours-full achieves improvement over Ours-partial and Ours-no-HIER by a large margin, which demonstrates the benefits of *back-track* (higher precision) and hierarchical exemplars (more expressive models). Note that Ours-partial already outperforms Strong-PS (Pishchulin et al. 2013b) and PoseMachine (Ramakrishna et al. 2014).

Fig. 4 shows some testing examples, which are selected with high diversity in poses. We can see that our method achieves accurate localization for most of the body joints, even in the case of large articulated deformation.

Method	Torso	ULeg	LLeg	UArm	LArm	Head	Avg
Strong-PS	88.7	78.8	73.4	61.5	44.9	85.6	69.2
PoseMachine	88.1	78.9	73.4	62.3	39.1	80.9	67.6
IDPR	92.7	82.9	77.0	69.2	55.4	87.8	75.0
Ours-partial	89.2	79.5	73.6	65.8	50.3	85.6	71.3
Ours-no-HIER	85.4	75.3	66.7	54.9	37.5	82.5	63.7
Ours-full	93.5	84.4	78.3	71.4	55.2	88.6	76.1
Ours-full (PC)	93.7	82.2	76.0	68.6	53.2	88.3	74.2

Table 1: Comparison of pose estimation results (PCP) on LSP dataset. Our method achieves the best performance.

LSP Extended Dataset (PC Annotations)

To have fair comparisons with (Toshev and Szegedy 2014; Tompson et al. 2014), we train and test our models on LSP extended dataset using PC annotations. Altogether, we have 11,000 training images and 1000 testing images. As the quality of the annotations for the additional training images varies a lot, we manually filter out about 20% of them. We also augment the training data through flipping and rotation.

We use Percentage of Detected Joints (PDJ) to evaluate the performance, which provides an informative view of the localization precision. In this experiment, we evaluate the baseline of our method (referred to as “Ours-base”) by only using the first term in Eq. 1. It is equivalent to localizing the parts independently. In Fig. 5, we plot the detection rate vs. normalized error curves for different methods. From the curves, we can see that Ours-base already achieves better accuracy than (Toshev and Szegedy 2014) except for Knee. It demonstrates that a detector that scores the part appearance is more effective than a regressor that predicts the part offset. Ours-full achieves significant improvement over Ours-base by incorporating the multi-level spatial models. Our method is also comparable to (Tompson et al. 2014) which enjoys the benefit of jointly learning appearance models and spatial context. (Tompson et al. 2014) has higher accuracy on the lower arms, while we have better results on the lower legs. Also note that (Tompson et al. 2014) requires delicate implementation of a sophisticated network architecture, while our method allows the use of off-the-shelf DCNN models.

CUB-200-2011 Bird Dataset

We also evaluate our method on the CUB-200-2011 bird dataset, which contains 5,994 images for training and 5,794 images for testing. Each image is annotated with image locations for 15 parts. We also augment the training data through flipping and rotation. As birds are less articulated than humans, we design a three-level hierarchy for birds. The first level contains the atomic parts; the second level has three composite parts (Head, Belly&Legs, and Back&Tail); and the third level corresponds to the whole bird. Although we did not prove that the manually-designed hierarchy is optimal, we empirically find that it facilitates the prediction of part relations.

We use PCP to measure performance. In the bird dataset, a correct part detection should be within 1.5 standard deviation of an MTurk worker’s click from the ground-truth location. For a semi-rigid object such as bird with sufficient

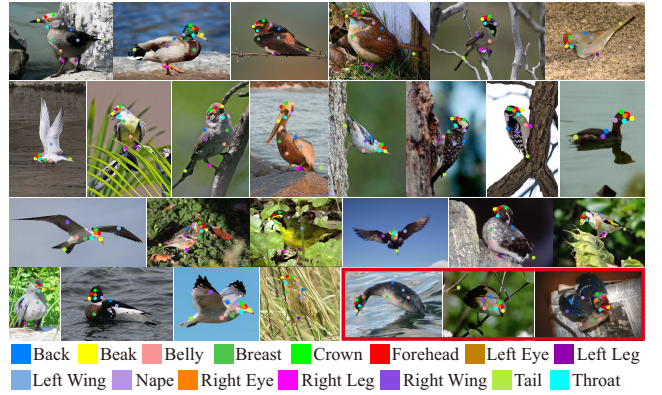


Figure 6: Qualitative results of part localization on CUB-200-2011 bird dataset. Failures are denoted with red frames, where some parts dramatically deviate from the correct locations due to large deformation and noisy patterns. The color codes are shown at the bottom.

training samples, directly applying exemplar-based models can produce very good results. Therefore, we replace the part detectors in (Liu and Belhumeur 2013) with DCNN-based detectors (only targeting the atomic parts), which will be referred to as “DCNN-CoE”.

We compare the results of different methods in Tab. 2, including CoE (Liu and Belhumeur 2013) and Part-pair (Liu, Li, and Belhumeur 2014). First, DCNN-CoE outperforms CoE significantly, demonstrating that DCNN is much more powerful than the conventional classification model (e.g., SVM). DCNN-CoE also outperforms Part-pair with much less overhead, thanks to the efficiency of multi-class detector. Using our new method, the localization accuracy is further improved. Ours-partial improves slightly over DCNN-CoE, which is reasonable as Ours-partial is essentially multi-level DCNNs plus multi-level exemplars, and the flexibility from our multi-level exemplars has limited effect for semi-rigid objects. Also note that Ours-partial uses an incomplete scoring function. By considering the full scoring function, Ours-full achieves the best results on all the parts.

Some qualitative results are shown in Fig. 6. From the examples, we can see that our method is capable of capturing a wide range of poses, shapes and viewpoints. In addition, our method localizes the bird parts with very high precision.

Conclusion

In this paper, we propose a novel approach for articulated pose estimation. The approach exploits the part relations at different levels of granularity through multi-scale DCNN-based models and hierarchical exemplar-based models. By incorporating DCNN-based appearance models in the spatial terms, our method couples spatial models with them, thus better adapting to the particular test image than otherwise. By introducing hierarchy in the exemplar-based models, we enjoy much more expressive spatial models even if the training data are limited. In addition, We propose an efficient algorithm to infer “good-enough” part configura-

Method	Ba	Bk	Be	Br	Cr	Fh	Le	Ll	Lw	Na	Re	Rl	Rw	Ta	Th	Total
CoE	62.1	49.0	69.0	67.0	72.9	58.5	55.8	40.9	71.6	70.8	55.5	40.5	71.6	40.2	70.8	59.7
Part-pair	64.5	61.2	71.7	70.5	76.8	72.0	69.8	45.0	74.3	79.3	70.1	44.9	74.4	46.2	80.0	66.7
DCNN-CoE	64.7	63.1	74.2	71.6	76.3	72.9	69.0	48.2	72.6	82.0	69.2	47.9	72.3	46.8	81.5	67.5
Ours-partial	65.1	64.2	74.6	72.4	77.1	73.8	70.2	48.4	73.2	82.5	70.6	48.7	73.0	48.3	82.2	68.3
Ours-full	67.3	65.6	75.9	74.4	78.8	75.3	72.5	50.9	75.4	84.7	72.8	50.4	75.2	49.9	84.2	70.2

Table 2: Comparison of part localization results on the CUB-200-2011 bird dataset. Our method outperforms the previous methods by a large margin. From left to right, the parts are: Back, Beak, Belly, Breast, Crown, Forehead, Left Eye, Left Leg, Left Wing, Nape, Right Eye, Right Leg, Right Wing, Tail, Throat, and Total.

tions from a less simplified formulation. These efforts together enable us to achieve state-of-the-art results on different datasets, which demonstrates the effectiveness and generalization ability of our method.

Acknowledgments

This work was supported by NSF awards 0968546 and 1116631, ONR award N00014-08-1-0638, and Gordon and Betty Moore Foundation grant 2987.

References

- [Belhumeur et al. 2011] Belhumeur, P. N.; Jacobs, D. W.; Kriegman, D. J.; and Kumar, N. 2011. Localizing parts of faces using a consensus of exemplars. *Proc. CVPR*.
- [Bourdev and Malik 2009] Bourdev, L., and Malik, J. 2009. Poselets: Body part detectors trained using 3d human pose annotations. *Proc. ICCV*.
- [Bourdev et al. 2010] Bourdev, L.; Maji, S.; Brox, T.; and Malik, J. 2010. Detecting people using mutually consistent poselet activations. *Proc. ECCV*.
- [Branson, Perona, and Belongie 2011] Branson, S.; Perona, P.; and Belongie, S. 2011. Strong supervision from weak annotation: Interactive training of deformable part models. *Proc. ICCV*.
- [Chen and Yuille 2014] Chen, X., and Yuille, A. 2014. Articulated pose estimation by a graphical model with image dependent pairwise relations. *Proc. NIPS*.
- [Dalal and Triggs 2005] Dalal, N., and Triggs, B. 2005. Histograms of oriented gradients for human detection. *Proc. CVPR* 1:886–893.
- [Felzenszwalb and Huttenlocher 2005] Felzenszwalb, P. F., and Huttenlocher, D. P. 2005. Pictorial structures for object recognition. *IJCV* 61(1):55–79.
- [Girshick et al. 2014] Girshick, R.; Donahue, J.; Darrell, T.; and Malik, J. 2014. Rich feature hierarchies for accurate object detection and semantic segmentation. *Proc. CVPR*.
- [Jain et al. 2014] Jain, A.; Tompson, J.; Andriluka, M.; Taylor, G. W.; and Bregler, C. 2014. Learning human pose estimation features with convolutional networks. *Proc. ICLR*.
- [Jia et al. 2014] Jia, Y.; Shelhamer, E.; Donahue, J.; Karayev, S.; Long, J.; Girshick, R.; Guadarrama, S.; and Darrell, T. 2014. Caffe: Convolutional architecture for fast feature embedding. *Proc. MM*.
- [Johnson and Everingham 2010] Johnson, S., and Everingham, M. 2010. Clustered pose and nonlinear appearance models for human pose estimation. *Proc. BMVC*.
- [Liu and Belhumeur 2013] Liu, J., and Belhumeur, P. N. 2013. Bird part localization using exemplar-based models with enforced pose and subcategory consistency. *Proc. ICCV*.
- [Liu, Li, and Belhumeur 2014] Liu, J.; Li, Y.; and Belhumeur, P. N. 2014. Part-pair representation for part localization. *Proc. ECCV*.
- [Pishchulin et al. 2013a] Pishchulin, L.; Andriluka, M.; Gehler, P.; and Schiele, B. 2013a. Poselet conditioned pictorial structures. *Proc. CVPR*.
- [Pishchulin et al. 2013b] Pishchulin, L.; Andriluka, M.; Gehler, P.; and Schiele, B. 2013b. Strong appearance and expressive spatial models for human pose estimation. *Proc. ICCV*.
- [Ramakrishna et al. 2014] Ramakrishna, V.; Munoz, D.; Hebert, M.; Bagnell, J. A.; and Sheikh, Y. 2014. Pose machines: Articulated pose estimation via inference machines. *Proc. ECCV*.
- [Rothrock, Park, and Zhu 2013] Rothrock, B.; Park, S.; and Zhu, S.-C. 2013. Integrating grammar and segmentation for human pose estimation. *Proc. CVPR*.
- [Sapp and Taskar 2013] Sapp, B., and Taskar, B. 2013. Modec: Multimodal decomposable models for human pose estimation. *Proc. CVPR*.
- [Sapp, Jordan, and Taskar 2010] Sapp, B.; Jordan, C.; and Taskar, B. 2010. Adaptive pose priors for pictorial structures. *Proc. CVPR*.
- [Sermanet et al. 2014] Sermanet, P.; Eigen, D.; Zhang, X.; Mathieu, M.; Fergus, R.; and LeCun, Y. 2014. Overfeat: Integrated recognition, localization and detection using convolutional networks. *Proc. ICLR*.
- [Sun and Savarese 2011] Sun, M., and Savarese, S. 2011. Articulated part-based model for joint object detection and pose estimation. *Proc. ICCV*.
- [Szegedy, Toshev, and Erhan 2013] Szegedy, C.; Toshev, A.; and Erhan, D. 2013. Deep neural networks for object detection. *Proc. NIPS*.
- [Tompson et al. 2014] Tompson, J.; Jain, A.; LeCun, Y.; and Bregler, C. 2014. Joint training of a convolutional network and a graphical model for human pose estimation. *Proc. NIPS*.

- [Toshev and Szegedy 2014] Toshev, A., and Szegedy, C. 2014. Deeppose: Human pose estimation via deep neural networks. *Proc. CVPR*.
- [Tsochantaridis et al. 2004] Tsochantaridis, I.; Hofmann, T.; Joachims, T.; and Altun, Y. 2004. Support vector machine learning for interdependent and structured output spaces. *Proc. ICML*.
- [Vondrick et al. 2013] Vondrick, C.; Khosla, A.; Malisiewicz, T.; and Torralba, A. 2013. Hoggles: Visualizing object detection features. *Proc. ICCV*.
- [Wang and Li 2013] Wang, F., and Li, Y. 2013. Learning visual symbols for parsing human poses in images. *Proc. IJCAI*.
- [Wang, Tran, and Liao 2011] Wang, Y.; Tran, D.; and Liao, Z. 2011. Learning hierarchical poselets for human parsing. *Proc. CVPR*.
- [Yang and Ramanan 2011] Yang, Y., and Ramanan, D. 2011. Articulated pose estimation with flexible mixtures-of-parts. *Proc. CVPR*.